twelfth which discusses explanation, all the remaining papers deal with this topic in a series of bite-sized chunks.

First is a discussion of a preference relation that can be used to construct higher order degrees of preference, such as the subtle notion that one can prefer the option of preferring not to smoke over the option of preferring to smoke, despite preferring to smoke. This is then followed by two closely related essays on the use of utility theory to make interpersonal comparisons, in other words, to provide a sensible means of choosing options that, as far as possible, satisfy everybody's preferences. Clearly, this is not straightforward, since it involves using some model of how to assess the well-being of a whole group from the well being of its individual members and the consequent dip in extremely murky waters. The final four essays are even more closely bound together, all being concerned with various aspects of the model of preferences introduced in the first part of Jeffrey's book *The Logic of Decision*. The first of these four, the thirteenth in the collection, summarizes the model in a few pages, and the second briefly compares it with other proposals, not surprisingly concluding that it is superior to them. The fifteenth essay then provides a detailed look at a system of axioms for the model of preference, and the sixteenth and final essay discusses at length the way in which preferences may be updated.

Reviewed by Simon Parsons, Advanced Computation Laboratory, Imperial Cancer Research Fund, London, UK

Inductive logic programming: techniques and applications by Nada Lavrac and Saso Dzeroski, Ellis Horwood, UK, 1993, pp 293, £39.95, ISBN 0-13-457870-8.

Inductive logic programming (ILP) is the fusion of symbolic machine learning with logic programming. Its aim is to learn logic programs from examples. The use of the logic programming formalism has advantages for machine learning. The formalism is elegant with a strong theoretical base, making it relatively easy to learn logic programs compared to, say, learning C programs. The formalism also makes it easier to incorporate background knowledge in the learning process, an essential to solve many learning problems. Logic programs (based on the first-order predicate calculus) are more expressive than traditional symbolic machine learning formalisms (based on the propositional calculus), such as decision trees. This means that a larger class of learning problems can be dealt with. The disadvantage of using such a powerful representation as ILP is that the search space is generally larger than for propositional learning; this is the reason that ILP has been, until recently, considered infeasible.

To appreciate the need for learning systems of the power of ILP, consider the following simple problem (from Quinlan, but very similar to many drug design problems). Consider a directional network of ten nodes, with each node connected to at most three others. In a traditional (propositional) machine learning formalism, an object must be specified by its values for a fixed set of attributes, and the rules learnt expressed as functions of these attributes. This would mean for our example that 30 (3×10) attributes would be needed, e.g. the attributes A1,B1,C1 would be the nodes to which node 1 is linked. Now consider the problem of expressing the concept "two nodes are linked to each other" using the propositional formalism. It gives a very complicated expression:

$(A1 = 2 \vee B1 = 2 \vee C1 = 2) \,\&\, (A2 = 1 \vee B2 = 1 \vee C2 = 1) \vee$
$(A1 = 3 \vee B1 = 3 \vee C1 = 3) \,\&\, (A2 = 3 \vee B2 = 3 \vee C2 = 3) \vee$
$(A1 = 4 \vee B1 = 4 \vee C1 = 4) \,\&\, (A2 = 4 \vee B2 = 4 \vee C2 = 4) \vee$
etc.

$(A2 = 3 \vee B2 = 3 \vee C2 = 3) \,\&\, (A3 = 2 \vee B3 = 2 \vee C3 = 2) \vee$
etc.

Now considering the problem in the ILP formalism, using only the predicate linked-to (X,Y) to express the fact that there is a directed link from X to Y. The concept "two nodes are linked to each other" can now be easily expressed by the PROLOG program:

linked-to-each-other (X, Y):- linked-to(X, Y), linked-to (Y, X).

This book is intended as an introduction to ILP. It is aimed at a reader with a basic knowledge of logic programming who wishes to learn the basics of new field of ILP, e.g. a knowledge engineer interested in the automatic synthesis of knowledge bases for expert systems. As such, it plays a useful role in gathering together information on ILP system and explaining them in a reasonably clear and concise way. The authors both have extensive experience of theoretical and experimental ILP. They originate from the AI Laboratory of the Jozef Stefan Institute (Ljubjana, Slovenia), which under its director Ivan Bratko, has had a long record of high quality work in both applied machine learning and logic programming.

The book is divided into two main parts: an introduction to the theory behind ILP systems, and a survey of the applications of ILP systems, the introduction being the most useful part. Much of the information given in this section is difficult to obtain from original sources (like too much of AI, many of the important papers have been published in hard to obtain conference proceedings and reports). The original technical papers are also in general quite difficult to read, as they assume a wide knowledge of machine learning and logic programming theory. The book describes in detail the basic ILP techniques of: relative least general generalization (RLGG), inverse resolution, search of refinement graphs, etc. Then proceeds to describe the best known ILP programs: MIS, FOIL, GOLEM, LINUS (developed at Ljubjana), etc.

The second part of the book describes examples of the use of ILP systems in practise. This part of the book I found less interesting, and I soon grew tired of reading page after page of experimental description. I would have been more convinced by it all if there was clearer evidence that the ILP systems were succeeding where traditional propositional systems had failed. I suggest that more work needs to be done in comparing ILP and existing propositional systems. I was also unconvinced by the extensive treatment given to the m-estimate method of measuring model fit—better methods now exist.

As a small quibble. Why is there a rather blurred picture of a strand of DNA on the cover of the book? I could find no reference to DNA in the text, the nearest being the description of an application of GOLEM to the problem of predicting a protein's secondary structure. Could not a more appropriate picture be found?

In conclusion, is ILP a chimera or the future of machine learning and knowledge acquisition? It is too early to tell, but this book provides the best available general introduction to the subject, and should be read by everyone interested in this exciting new area.

Reviewed by Ross D. King, Biomolecular Modelling Laboratory, Imperial Cancer Research Fund, London, UK

The phenomenon of commonsense reasoning: nonmonotonicity, action and information by Dimitrios Thanassas, Ellis Horwood, UK, 1992, pp 254, £37.95, ISBN 0-13-663634-9.

This book is, or claims to be, about commonsense reasoning. An example of commonsense reasoning that I feel has some relevance to this review is the following: if having read half of a certain book one has found nothing to lead one to believe that the book is worth reading, then it is reasonable not to bother reading the other half. I have to confess that I find myself in rather this position with respect to the book currently under review. I have read the first six chapters, and have not found the stomach to read the remaining four. If on that account the following review is discredited on the grounds that the reviewer could not be bothered to read the book through to the end, then so be it; nonetheless, I believe that what I have to say about the first half is quite likely to be of relevance to the second half also.

What, then, do we find in the first six chapters? The first chapter, a general critique of the current state of artificial intelligence (AI), need not detain us. The second is a broad survey of theories of